

Algorithmic Discrimination of Scientific Preprints

Hazards, Legal Implications, and Epistemic Consequences
of AI-driven Gatekeeping on Preprint Platforms

*An Interdisciplinary Analysis at the Intersection of
Science Theory, Digital Law, and AI Ethics*

Sven Mildner¹

¹Germania Magna Research Project, Briesnitzer Höhe 28, 01157 Dresden, Germany

Correspondence: germania.magna@smi-web.de

ORCID: <https://orcid.org/0009-0005-8248-4866>

doi: [10.5281/zenodo.15192847](https://doi.org/10.5281/zenodo.15192847)

Keywords: Preprint server, peer review, AI bias, freedom of science, Digital Markets Act, Epistemic Injustice, LLM hallucination, Open Science, § 824 BGB, equal treatment principles, algorithmic gatekeeping, epistemological diversity

Legal Notice

This paper serves scientific and informational purposes only. It does not constitute legal advice within the meaning of the German Legal Services Act (RDG). For specific legal action, consultation with a lawyer specialising in IT and media law is recommended.

Contents

1	Abstract	3
2	Introduction: The Preprint as a Legitimate Instrument of the Scientific Process	3
3	The Problem: LLMs as Distorting Quality Guardians	4
3.1	Statistical Correlation as an Epistemic Instrument	4
3.2	The Critical Misconstruction: Status vs. Outcome	4
3.3	Scenarios of Miscommunication in Practice	5
4	Legal Analysis: Implications for Researchers and Platform Operators	6
4.1	False Statement of Fact and Credit Endangerment (§ 824 BGB)	6
4.2	Competition Law: Disparagement under § 4 No. 1 UWG	6
4.3	Digital Markets Act (DMA): The Prohibition of Self-Preferencing	6
4.4	Digital Services Act (DSA): Transparency Obligations for Recommender Systems	7
4.5	Competition Law: § 19 GWB and Art. 102 TFEU	7
4.6	Freedom of Science: Art. 5 Para. 3 basic Law (GG)	7
5	Equal Treatment Principles and the Prohibition of Discrimination	7
5.1	Structural Unequal Treatment Based on Publication Status	7
5.2	The Question of Indirect Discrimination	8
6	Science-Theoretical Implications:	
	The Kuhn Dilemma in the Age of AI	8
6.1	Paradigm Shifts as Victims of Algorithmic Conformity	8
6.2	Epistemic Injustice (Miranda Fricker)	8
6.3	Erosion of Epistemic Diversity	9
7	Consequences for Operators of Preprint Servers	9
7.1	Reputational Damage in the Digital Ecosystem	9
7.2	Recommendations for Action for Platform Operators	9
7.2.1	Legal Measures	9
7.2.2	Communicative and Technical Measures	10
8	Responsibility of AI Developers: Structural Ignorance and Epistemic Duties	10
8.1	The Unintended Gatekeeper	10
8.2	The Epistemic Responsibility	10

8.3 The Principle of Hallucination through Statistical Correlation 11

9 Discussion: Limits of the Critique and Counter-Arguments 11

10 Conclusions and Recommendations 12

1. Abstract

The growing integration of generative AI systems (Large Language Models, LLMs) into scientific search and evaluation processes harbours a hitherto largely underestimated systemic risk: the algorithmic devaluation of preprints through linguistically erroneous or tendentious formulations. This paper examines how LLMs, through statistical pattern recognition, tend to incorrectly link the formal publication status of a work with its content quality. This mechanism leads to a *Mischaracterization* that may, in individual cases, fulfil the elements of a false statement of fact within the meaning of § 824 BGB and systemically endangers free competition of scientific ideas within the meaning of the *Digital Markets Act* (DMA) and the *Digital Services Act* (DSA).

It is shown that preprint platforms, as infrastructure of the Open Science movement, must be actively protected, since their algorithmic discreditation contradicts both equal treatment principles and the freedom of science guaranteed by Art. 5 para. 3 of the German Basic Law (GG). Recommendations for action are formulated for platform operators, researchers, and regulatory authorities.

Keywords: Preprint server, peer review, AI bias, freedom of science, Digital Markets Act, Epistemic Injustice, LLM hallucination, Open Science, § 824 BGB

2. Introduction: The Preprint as a Legitimate Instrument of the Scientific Process

The history of science is a history of incomplete, revisable, and contestable knowledge. Scientific knowledge acquisition does not proceed as a linear process from idea to secured truth, but as an iterative, social process of negotiation (Kuhn, 1962; Popper, 1935). Preprints are an established and functionally indispensable instrument that allows scientists to feed ideas, models, and theses transparently into the scientific community before a formal review procedure has been completed or even initiated.

Preprint servers such as EarthArXiv, arXiv, bioRxiv, SSRN, or ResearchGate fulfil several key functions in the scientific system:

1. **Acceleration of discourse:** Complex, interdisciplinary, and iterative models can be made accessible to the scientific community in a timely manner, without years-long waiting times in the editorial processes of established journals.
2. **Prevention of structural barriers:** Paradigm shifts that radically question existing disciplines or completely re-link subject areas encounter structural resistance in the classical, often conservative peer-review process of established journals — a mechanism described by Kuhn (1962) as a property of “normal science.”

3. **Democratisation of knowledge:** Preprints enable barrier-free, cost-free access to scientific output, independent of commercial gatekeeper publishers (Willinsky, 2006).
4. **Transparency of process:** The open documentation of intermediate stages, revisions, and extensions of a model is an element of good scientific practice and corresponds to the spirit of reproducible research.

The use of a preprint server is therefore not a deviation from scientific norms — it is a deliberate, methodically justified decision that possesses full legitimacy in the scientific context.

3. The Problem: LLMs as Distorting Quality Guardians

3.1 Statistical Correlation as an Epistemic Instrument

Large Language Models (LLMs) such as Gemini, GPT-4, or comparable systems are trained on enormous text corpora representing the published knowledge of humanity. In doing so, they do not learn logical relationships in the strict philosophical sense, but statistical co-occurrence patterns: which concepts appear in which linguistic neighbourhoods?

This leads to a structural problem that is decisive for the present context. In the text corpus on which LLMs are trained, the term “*not peer-reviewed*” appears statistically disproportionately frequently in the vicinity of expressions such as “*not verified*,” “*quality deficit*,” “*not credible*,” or “*failed*.” This does not, however, reflect the reality of the scientific process, but merely a medial and discursive bias of the underlying training corpus.

Formally, this relationship can be described as follows. Let $P(\text{“qualitatively inferior”} \mid \text{“not peer-reviewed”})$ be the conditional probability with which an LLM generates the term *qualitatively inferior* in the vicinity of *not peer-reviewed*. Due to the statistical distribution in the training corpus, the following may hold:

$$P(\text{“qualitatively inferior”} \mid \text{“not peer-reviewed”}) \gg P(\text{“under discussion”} \mid \text{“not peer-reviewed”}) \quad (1)$$

Although the latter represents the correct epistemic classification, the former is linguistically preferred because it occurs more frequently in the training data.

3.2 The Critical Misconstruction: Status vs. Outcome

The most serious manifestation of this distortion is a specific semantic confusion that may be termed **Mischaracterization of Process**: an LLM-generated system does not

describe a preprint as a work that *is in the stage before or outside* a review procedure, but implies that it has *failed* such a procedure.

This distinction is fundamental both legally and in terms of philosophy of science:

Table 1: Comparison of correct and erroneous AI classification of preprints

Category	Correct Classification	Erroneous LLM Formulation
Procedural	“Is in open discourse”	“Has failed the controls”
Temporal	“Not yet peer-reviewed”	“Was rejected”
Legal	Value judgement (permissible)	False statement of fact (problematic)
Epistemic	Ongoing process → discussion → revision	Binary final verdict: passed / failed

The formulation “*has not passed the strict review processes of recognised specialist journals*” describes a historical event — namely an active rejection by a journal — that by definition has not occurred in the case of a preprint. It is therefore not a value judgement but a factually false statement.

3.3 Scenarios of Miscommunication in Practice

The following anonymised case example illustrates the mechanism:

Anonymised Case Example

A researcher publishes a methodologically complex, interdisciplinary model on a recognised preprint server. The model is continuously being developed and fundamentally challenges established doctrines in its field. The author has deliberately not submitted the work for review at a journal, as the model is iterative in nature and open scientific discourse on a preprint platform is more appropriate to this stage of development.

An LLM-powered system, responding to a query about this researcher, formulates:

“The works were predominantly uploaded as preprints and have not passed the strict review processes of recognised specialist journals.”

This statement is incorrect and potentially damaging in several respects: it implies an active rejection by colleagues that demonstrably did not take place; it devalues the infrastructure of the preprint server with the same judgement; and it is generated by a system perceived by millions of users as an authoritative source of information.

4. Legal Analysis: Implications for Researchers and Platform Operators

4.1 False Statement of Fact and Credit Endangerment (§ 824 BGB)

§ 824 para. 1 BGB provides:

“Anyone who, contrary to the truth, asserts or disseminates a fact that is likely to endanger the creditworthiness of another [...] shall be liable to the other party to compensate for the damage resulting therefrom.”

The statement that a preprint has “*not passed*” the review processes of a journal potentially fulfils all elements of the offence under prevailing civil law doctrine:

- **Statement of fact:** The formulation “not passed” designates a specific historical event (a rejection) that is amenable to verification — as opposed to a mere value judgement.
- **Untrue fact:** The event (rejection by a journal) has by definition not taken place if the work was never submitted.
- **Credit endangerment:** Such a statement can damage the scientific reputation of the author and — if disseminated systemically — the reputation of the hosting preprint platform.

Legal Classification

The question of liability for AI-generated content has not yet been conclusively resolved in law. Under the Product Liability Act and the EU’s emerging AI liability rules (European Commission, 2022), it is emerging that the operator of the AI system may be held accountable for damaging outputs if insufficient care was applied in system design.

4.2 Competition Law: Disparagement under § 4 No. 1 UWG

The Act against Unfair Competition (UWG) also protects platform operators against deliberate disparagement. § 4 No. 1 UWG prohibits actions likely to damage a competitor or its services by making untrue statements or defaming them. While there is no classic competitive relationship in the narrow sense between an AI provider and a preprint server, both act as commercial actors in the market for information infrastructure, which makes an extended application of UWG norms possible.

4.3 Digital Markets Act (DMA): The Prohibition of Self-Preferencing

Since March 2024, the EU’s *Digital Markets Act* (DMA) has been in force and obliges so-called “gatekeepers” — very large online platforms — to strict neutrality towards third-party services (Art. 6 para. 5 DMA) (European Parliament and Council, 2022a). If an AI system integrated into a dominant search engine systematically devalues Open

Access preprint servers while presenting proprietary or favoured information services more positively, this raises the likelihood of a violation of the DMA's prohibition of discrimination.

4.4 Digital Services Act (DSA): Transparency Obligations for Recommender Systems

Art. 27 DSA obliges platforms to disclose the parameters that determine which information is presented in which order (European Parliament and Council, 2022b). If an AI-powered recommender system systematically associates preprint content with negative quality judgements without the underlying criteria being transparent to users, this contradicts the spirit of the DSA.

4.5 Competition Law: § 19 GWB and Art. 102 TFEU

Should a dominant search engine (market share above 80 % in the German search engine market) disadvantage preprint platforms without objective justification through its integrated AI system, competition law applies. § 19 para. 1 GWB (abuse of a dominant position) and Art. 102 TFEU (abuse of a dominant position in the internal market) provide the legal framework for competition law complaints to the Federal Cartel Office or the EU Commission (Bundeskartellamt, 2019).

4.6 Freedom of Science: Art. 5 Para. 3 Basic Law (GG)

Art. 5 para. 3 GG guarantees that “art and science, research and teaching [...] shall be free.” This fundamental right protects not only the freedom to conduct research, but also the freedom to disseminate scientific findings and present them for discussion. If AI systems, functioning as de-facto information infrastructure of society, algorithmically discredit certain forms of scientific communication (here: preprints), an indirect endangerment of fundamental rights arises.

5. Equal Treatment Principles and the Prohibition of Discrimination

5.1 Structural Unequal Treatment Based on Publication Status

Equal treatment requirements are found at various levels of German and European law. In the context of scientific communication, it is relevant that unequal treatment of scientists solely on the basis of the formal publication status of their work — not its content — constitutes a differentiation that is not objectively justified.

The scientific value of a work is a question of content that must be judged by a scientific community. The decision to first publish a work as a preprint is a strategic and methodically justified decision that:

- may correspond to the early stage of an iterative model,

- reflects the desire for open scientific feedback,
- responds to the structural limitations of classical peer-review procedures for interdisciplinary or revolutionary theses, or
- is simply a question of resources and time horizon.

A blanket devaluation on the basis of publication status is incompatible with the equal treatment principle, since identical subject matters (scientific works) are treated differently without objective justification.

5.2 The Question of Indirect Discrimination

Particularly for interdisciplinary or system-critical research approaches, which structurally can less often be incorporated into commercial journal systems, the algorithmic devaluation of preprints acts as indirect discrimination: it does not affect all researchers equally, but disproportionately those whose work is methodologically, thematically, or institutionally located outside the academic mainstream.

6. Science-Theoretical Implications:

The Kuhn Dilemma in the Age of AI

6.1 Paradigm Shifts as Victims of Algorithmic Conformity

Thomas S. Kuhn showed in his history of science that genuine scientific revolutions always began as outsider opinions initially marginalised by “normal science” (Kuhn, 1962). Copernicus, Galileo, Semmelweis, Wegener — all were rejected by the established institutions of their time before their theories became the new paradigm.

In the age of AI-powered information systems, a new form of this mechanism emerges: not just institutions, but algorithmic systems assume the gatekeeping function. Since these systems are statistically trained on the existing knowledge canon, they structurally replicate the intellectual status quo. They are formally predisposed to marginalise paradigm shifts — not through malice, but through the inherent logic of their statistical architecture.

6.2 Epistemic Injustice (Miranda Fricker)

The philosopher Miranda Fricker developed the concept of **epistemic injustice**: the injustice that befalls an individual or group when their ability to produce, communicate, or be recognised for knowledge is systematically undermined (Fricker, 2007). Fricker distinguishes:

- **Testimonial injustice**: A knowledge contribution is not taken seriously because epistemic trust is denied to the originator on the basis of an identity characteristic.
- **Hermeneutical injustice**: An individual or group suffers from a deficit of interpre-

tive resources to name and communicate their own experience.

In the context of algorithmic preprint discrimination, one may speak of an **algorithmic epistemic injustice**: the AI system withdraws epistemic trust from the originator not on the basis of a personal characteristic, but on the basis of a publication characteristic — with identical effect on their scientific reception.

6.3 Erosion of Epistemic Diversity

If AI systems are used pervasively as the primary source of information and structurally devalue Open Science publications, the cultural and scientific corridor of what counts as valid knowledge narrows. This has long-term consequences:

- Researchers are placed under pressure to choose the expensive and often lengthy route of commercial journal publications to avoid algorithmic devaluation.
- The **democratisation of knowledge** sought by the Open Science movement is counteracted.
- Social and political decision-makers who rely on AI-generated summaries receive a distorted picture of scientific discourse.

7. Consequences for Operators of Preprint Servers

7.1 Reputational Damage in the Digital Ecosystem

Preprint servers depend on the trust of the scientific community. This trust feeds on the correct perception of their function as open, transparent discussion platforms — not as repositories for rejected or inferior works. If AI systems, used daily by millions of people, structurally paint a negative picture of preprint platforms, a **cumulative reputational damage** arises that may manifest in measurable economic consequences:

- Decline in submission numbers by researchers fearing reputational harm,
- Reduction of funding and institutional partnerships, and
- Loss of user trust and visibility in the information ecosystem.

7.2 Recommendations for Action for Platform Operators

7.2.1 Legal Measures

1. **Documentation:** Systematic capture and archiving of AI-generated text passages that incorrectly or disparagingly represent preprint content or the platform itself — over a defined observation period.
2. **Formal warning (Abmahnung):** Formal cease-and-desist letter to the AI operator demanding cessation of demonstrably false statements of fact (§ 824 BGB, §§ 3, 4 UWG).

3. **Competition complaint:** Filing a complaint with the Federal Cartel Office or the EU Commission if a dominant operator is involved in the discrimination.
4. **DSA complaint:** Initiation of a formal transparency request under Art. 27 DSA with the competent national coordinators for digital services.
5. **Civil lawsuit:** In the event of demonstrable economic damage, action for injunctive relief and damages under §§ 823, 824 BGB.

7.2.2 Communicative and Technical Measures

1. **Active framing:** Strategic communication towards AI providers and the public that clearly presents the function and value of preprint servers and can act as a corrective to misinformation.
2. **Metadata optimisation:** Technical optimisation of machine-readable metadata in preprint entries to facilitate correct classification of publication status by algorithmic systems.
3. **Institutional cooperation:** Collaboration with other preprint platforms, scientific organisations, and Open Science initiatives to collectively respond to misrepresentations and demand guidelines for AI operators.
4. **Advocacy with AI providers:** Direct contact with the developers of the relevant AI systems to point out the bias and request integration of protective measures (*guardrails*) for the special status of preprints.

8. Responsibility of AI Developers: Structural Ignorance and Epistemic Duties

8.1 The Unintended Gatekeeper

It would be erroneous to understand the described mechanism solely as malicious, deliberate discrimination. In the majority of cases, it is the result of **naïve scaling**: the developers of large language models train their systems on the largest possible quantity of text and focus on global performance metrics — without sufficient sensitisation to the subtle institutional rules of subsystems such as the Open Science infrastructure.

The outcome is nevertheless identical to active gatekeeping: the system marginalises unconventional approaches algorithmically, not through a deliberate decision, but through the absence of an explicit counterstrategy (missing *guardrails*).

8.2 The Epistemic Responsibility

LLMs are increasingly assuming the function of a societal knowledge infrastructure. With this function comes an **epistemic responsibility** that goes beyond classic product

liability questions. AI developers bear the duty to:

- document and communicate the inherent biases of their systems,
- implement specific protective measures for vulnerable knowledge domains (such as Open Science), and
- always make transparent to users the context-dependence and error-proneness of AI-generated quality judgements.

8.3 The Principle of Hallucination through Statistical Correlation

A central technical problem lies in the mechanism of **Association Learning**: the language model generates text not through logical inference, but by calculating the most probable next token sequence in a given context. The confusion of “*not yet peer-reviewed*” with “*has not passed peer review*” is therefore not the result of a false logical inference, but the consequence of a strong statistical association in the training corpus.

Formally, for the token generation of an LLM with parameters Θ , given context c , and generated token t_i :

$$t_i = \arg \max_t P_{\Theta}(t \mid t_1, t_2, \dots, t_{i-1}, c) \quad (2)$$

This formula illustrates that the output is entirely determined by the probability distribution P_{Θ} approximated during training — and thus by the biases of the training corpus. Legal liability therefore shifts entirely to the developers and operators (European Commission, 2022).

9. Discussion: Limits of the Critique and Counter-Arguments

The present analysis should not be misunderstood as a blanket critique of AI systems or as an argument against any form of quality evaluation. It is equally permissible and scientifically necessary to identify content-related weaknesses in a preprint — methodological errors, logical inconsistencies, insufficient data basis. What is legally and ethically problematic is solely the linking of publication status with a negative quality judgement as a blanket conclusion.

Moreover, it must be acknowledged that peer review remains a valuable instrument of quality assurance — not as an exclusion mechanism, but as one of several possible paths of scientific verification. The critique is directed not against the peer-review system, but against its algorithmic hypostatisation as the *sole* measure of scientific legitimacy.

Important Note

The algorithmic devaluation of preprints is not an isolated technical problem, but the expression of a deeper societal bias encoded in the training corpora of the models. Technical solutions (guardrails, curated fine-tuning) are necessary but not sufficient: they must be complemented by regulatory frameworks and a critical digital science culture.

10. Conclusions and Recommendations

The foregoing analysis shows that the algorithmic devaluation of preprints by AI language models constitutes a multifaceted problem combining science-theoretical, legal, and societal dimensions. The following conclusions may be drawn:

1. **Legally**, specific LLM formulations that equate preprint status with failure in review procedures constitute potential cases of unfair false statements of fact that may give rise to liability under § 824 BGB, §§ 3, 4 UWG, DMA, and DSA.
2. **From the perspective of philosophy of science**, the algorithmic hierarchisation of publication forms contradicts the processual character of scientific knowledge and endangers the epistemic diversity from which scientific revolutions emerge (Kuhn, 1962).
3. **For preprint operators**, concrete options for action exist: from systematic documentation via civil-law cease-and-desist letters through to competition and DSA complaints with the competent authorities.
4. **For AI developers**, a clear ethical and increasingly also legal duty follows to implement specific protective measures for Open Science infrastructure and to explicitly communicate to users the distinction between publication status and content quality (European Commission, 2022).
5. **At the societal level**, the growing dependence on AI-generated information requires critical reflection on the epistemic authority attributed to these systems — and a collective effort to defend the democratisation of knowledge against algorithmic gatekeeping tendencies (Fricker, 2007; Willinsky, 2006).

Concluding Remark

For scientific work, this case provides a precise example: it shows that the problem of Big-Tech AI models lies not only in deliberate censorship, but in a **structural ignorance** towards agile, democratic, and open science structures. The system inherently tends to reinforce monopolies (such as the classical peer-review publishing system) and to inadvertently sabotage the democratisation of knowledge through Open Science. When AI systems take over the interpretive authority over what counts as “valid knowledge,” technology transforms from a research tool into an epistemic control instrument.

Legal Disclaimer

This paper serves scientific and informational purposes only. It does not constitute legal advice within the meaning of the German Legal Services Act (RDG) or equivalent national legislation. For specific legal steps, consultation with a lawyer specialising in IT and media law is strongly recommended.

References

- Bundeskartellamt (2019). *Misuse proceeding against Facebook*. Az. B6-22/16. Bonn: Federal Cartel Office.
- Bundesgerichtshof (2014). Judgment of 14 May 2013 — VI ZR 269/12 (Search engine results and personal rights). *BGH Decision*. Karlsruhe.
- European Commission (2022). *Proposal for a Directive of the European Parliament and of the Council on adapting non-contractual civil liability rules to artificial intelligence*. COM(2022) 496 final. Brussels.
- European Parliament and Council (2022a). *Regulation (EU) 2022/1925 on contestable and fair markets in the digital sector (Digital Markets Act)*. *Official Journal of the EU*, L 265. Brussels.
- European Parliament and Council (2022b). *Regulation (EU) 2022/2065 on a single market for digital services (Digital Services Act)*. *Official Journal of the EU*, L 277. Brussels.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Popper, K. R. (1935). *Logik der Forschung*. Wien: Springer.
- Willinsky, J. (2006). *The Access Principle: The Case for Open Access to Research and Scholarship*. Cambridge, MA: MIT Press.